

(12) UK Patent Application (19) GB (11) 2 344 029 (13) A

(43) Date of A Publication 24.05.2000

(21) Application No 9821555.1

(22) Date of Filing 02.10.1998

(71) Applicant(s)
General Datacomm Inc
(Incorporated in USA - Delaware)
Straits Turnpike, Route 63, Middlebury,
Connecticut 06762-1299, United States of America

(72) Inventor(s)
David Banes

(74) Agent and/or Address for Service
Mathys & Squire
100 Grays Inn Road, LONDON, WC1X 8AL,
United Kingdom

(51) INT CL⁷
H04L 12/56

(52) UK CL (Edition R)
H4P PPNB PPS
H4K KTK


(56) Documents Cited
GB 2187917 A GB 2166320 A WO 95/12265 A1
JP 580173932 A US 5251209 A

(58) Field of Search
UK CL (Edition R) H4K KTK , H4P PPNB PPS
INT CL⁷ H04L 12/56
Online Databases: WPI, EPODOC, JAPIO

(54) Abstract Title
Transmission of data packets of different size and priority

(57) The transmission of a lower priority data packet over a network is interrupted so that a data packet of a higher priority can be transmitted, after which the transmission of the lower priority data packet is resumed. A code is inserted into the lower priority packet signifying interruption. Interruption may be decline if a count of interruptions reaches a maximum count. Data packets have a length or type field in the header signifying long or short packets. The packet size is selected according to voice or data applications. Application is to ATM or Ethernet networks.


Fig. 5



GB 2 344 029 A

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

The print reflects an assignment of the application under the provisions of Section 30 of the Patents Act 1977.

*Fig. 1*





Fig. 2

*Fig. 3*


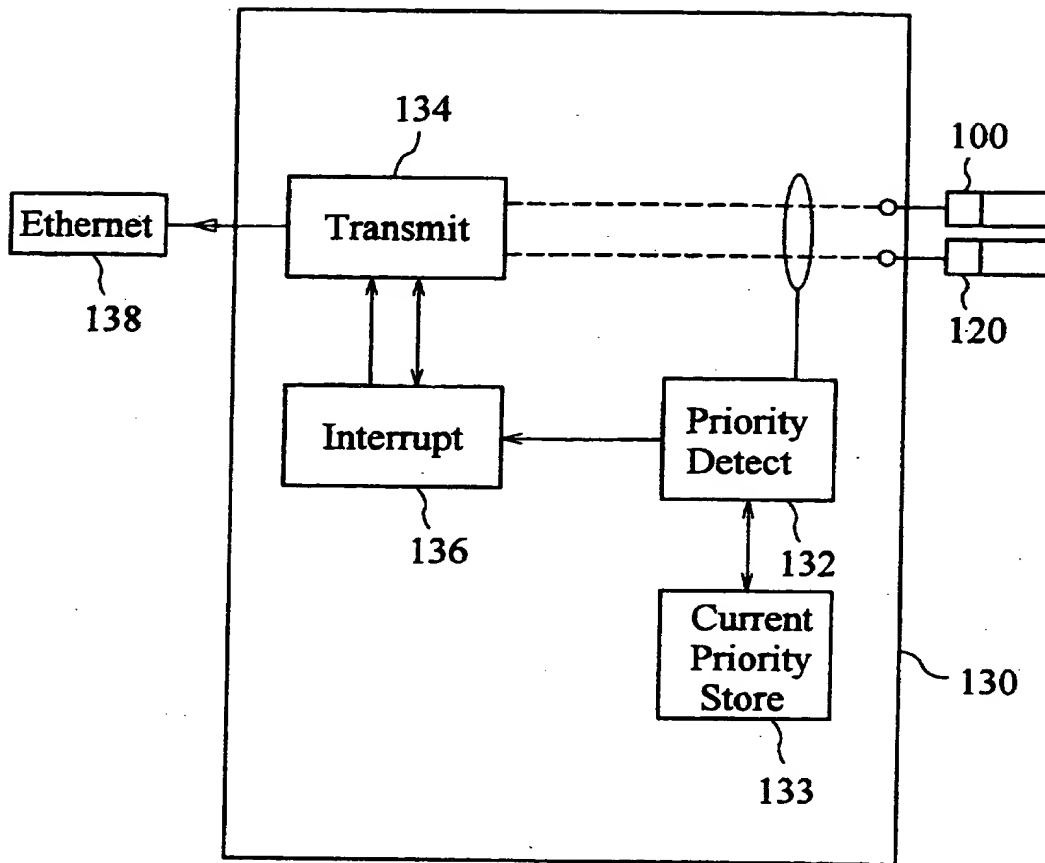
*Fig. 4*

Fig. 5

COMMUNICATION NETWORK

The present invention relates to a data communication network.

There are a number of data communication protocols in common usage. As is well-known, a data communication protocol is defined as a series of hierarchal layers, the lowest layer defining the physical characteristics of the transfer medium (for example cable or fibre optic), and higher levels defining software interfacing to applications.

One example of a known communications network is an Ethernet network. In such a network data is transported in packets of a variable length of up to approximately 1500 bytes. This type of network is well suited to local area network (LAN) applications for transferring computer data files or for transferring TCP/IP data, which is generally transferred in large blocks. It is not, however, as well suited to voice or to telecommunications networks in which data needs to be switched rapidly and frequently between destinations.

Another example of a network is an ATM network. In such a network, data is transmitted in fixed length cells of 53 bytes, 48 bytes of which contain the data or "payload" and 5 bytes of which contain a header. Because of the small size of the packets, packets can be readily switched between several different destinations, and the network is well suited to the needs of telecommunications and voice traffic. Larger blocks of data can be accommodated simply by transmitting them over a series of ATM packets. ATM networks have been successfully used to carry both voice and computer data. The small size of the packets enables efficient statistical multiplexing.

In both the above prior art systems, the data packets are of a fixed maximum length. Division of a long data block into separate packets is performed at a higher level, by software.

The inventor has appreciated that, whilst an ATM network can provide

an efficient and effective network for transmitting both voice and data, efficiency can be impaired when large amounts of data are transmitted. On the other hand, a network such as an ethernet may cause unacceptable delays between packets and cannot readily be used for voice traffic. Accordingly, it is an aim of the present invention to provide a network which can efficiently transfer both voice and computer data or which can transfer large blocks of data efficiently without unduly delaying transit of smaller or more important data packets.

In a first method aspect, the invention provides a method of communicating data over a network comprising transmitting a first data packet, receiving a higher priority data packet, interrupting transmitting the first data packet to transmit the higher priority data packet, and resuming transmission of the first data packet.

In this way, large blocks of data can be efficiently transmitted, without incurring the overhead associated with splitting them into several smaller packets, but without creating a bottleneck and holding up more urgent data.

Preferably interrupting transmission comprises sending a special code to signify interruption. For example, in the case of high speed Ethernet, 10 bit coding is used to carry 8 bit data, and an unused code can be used to flag an interruption.

Interruptions may be nested and a variety of priority levels may be provided. Preferably, however, only two priority levels are provided, as this reduces the risk of multiple interruptions and possible corruption of a long data packet.

A count may be kept of the number of interruptions of a lower priority packet, and further interruptions may be denied after a maximum count is reached, or the priority threshold required for interruption may be increased.

This may prevent long data packets from being unduly delayed by multiple interruptions.

5 In a related apparatus aspect, the invention provides apparatus for communicating data over a network comprising means for transmitting a first data packet, means for receiving a higher priority data packet, means for interrupting transmitting the first data packet to transmit the higher priority data packet, and means for resuming transmission of the first data packet.

10 In a second method aspect, the invention provides a method of communicating data over a network comprising selecting one of at least two predetermined packet sizes based on the data type and transmitting a packet of the selected size over the network.

15 In this way, rather than using fixed length packets, a packet size is "intelligently" selected to be appropriate to the data being transmitted. For example, voice data may be transmitted in a relatively short packet, preferably of approximately 64 bytes or less, and other data may be transmitted in a relatively long packet, preferably of at least about 256 bytes, one kilobyte, or more. Of course, such a method may require a more complex routing system and physical arrangement, which must deal with data packets of different sizes. Surprisingly, however, an improvement in network performance may be
20 gained by dividing data in this way prior to transmission.

25 The length of the data packet could be determined by routing and receiving components simply by monitoring data wave forms on the network. More preferably, however, the method includes providing at least one flag bit indicating the length of the data packet, preferably in a header associated with the data packet.

As well as differing in size, the internal format of the packets may differ. For example, short packets may be similar to ATM cells and long packets may

be similar to Ethernet frames. Both types of packet may be packaged and efficiently transported over a common medium, and readily separated for transmission over dedicated networks.

5 In a second apparatus aspect, the invention provides a data network including means for transmitting data of a first type as a series of packets of a first length over a communication medium and means for transmitting data of a second type as packets of a second length over the same communication medium.

10 The network preferably further includes routing means connected to the communication medium for routing packets of both said first length and said second length.

15 By providing only a relatively limited discrete number of packet lengths, the advantages of hardware routing available in conventional networks which employ a fixed single packet length may still be attained with only a slight increase in router complexity. For example, if two packet lengths are provided only a single bit is required in the header to identify the packet as long or short.

20 A plurality of discrete packet lengths may, however, be provided. Indeed, the packet length may be variable. However, with variable length data packets, some of the advantages of predicability of packet length and statistical multiplexing may be lost, and more complex routers, with a greater quantity of redundant buffer memory may be required.

25 In a third aspect, the invention provides a transmitter for connection to a communication network comprising means for selecting a data packet length from a plurality of predetermined packet lengths; means for providing a header for the data packet, the header including at least one flag bit indicating the length or format of the data packet; and means for transmitting the data packet over the communication medium.

In a fourth aspect, the invention provides a receiver for receiving data from a communication network, the receiver including means for reading a flag bit contained in the header of a data packet transmitted on the communication medium; means for determining the size or format of data packet to receive based on the information; and means for receiving a data packet of said size or format.

Embodiments of the invention will now be described by way of example with reference to the accompanying drawings in which:

Fig. 1 is a schematic diagram of data packets according to the second aspect of the invention;

Fig. 2 is a schematic diagram of a computer network embodying the invention;

Fig. 3 is a simplified schematic diagram of a switch for use in a network of Fig. 2;

Fig. 4 is a schematic diagram of transmission of data packets according to the first aspect of the invention; and

Fig. 5 is a schematic diagram of apparatus for implementing the transmission of Fig. 4.

An embodiment of the invention based on an ATM network will be described. Whilst this embodiment is based closely on ATM protocol, for ease of implementation by modification of existing ATM network components it is to be understood that the principles are applicable to other data formats.

Referring to Fig. 1 a data packet according to the invention comprises a header 10 and a data portion 20. The header comprises routing information 12 and a length or type field 14, which in this embodiment comprises a single bit signifying "long" or "short". The header may also contain information such as priority information.

The lengths of the short and long packets may be determined according to the intended application of the network. The short data packets are preferably adapted for voice communication and have a data portion of no more than about 64 bytes, preferably no more than about 48 bytes or even no more than about 32 or 16 bytes. The total length of a short packet including the header is preferably less than about 128 bytes. The long data packets are preferably significantly longer than the short data packets, preferably at least about 256 bytes, more preferably about 512, 1K 2K bytes or more. Advantageously, each long data packet is sufficiently long to contain a conventional disk block of a data file.

In this embodiment, the short data packets are very similar to the packets used in a conventional ATM network; this enables the embodiment to be implemented by modifying existing ATM network components. The flag bit may be anywhere in the packet, preferably in the header, and advantageously as early as possible in the header, to enable routing or identification of the packet to commence without substantial delay of the remainder of the packet. In this embodiment, the flag bit is a single bit appended to the front of the packet, and this enables packets to be directed either to a conventional ATM router or to a modified "long ATM" router arranged to handle ATM-type packets having a data content of 1024 bytes in a similar way to a conventional router.

A network as shown in Fig. 2 includes a "short" packet assembler/disassembler 40 for interfacing a voice communication device such as a telephone 42 to the network. The short PAD 40 may receive signals in either analog or digital form from the telephone 42 and produces ATM-type packets with an additional initial flag bit signifying that the packet is a short packet. A single PAD 40 may be coupled to a plurality of voice input devices; in such a case, the PAD 40 is arranged to perform final routing of received packets to the appropriate voice terminal.

In a similar manner, a long PAD 50 is arranged to receive data from the data input terminals such as a computer 52. This produces packets similar to conventional ATM packets, but having a considerably longer data field, and an initial flag bit indicating that the packet is long. Optionally, the packet may contain a cyclic redundancy check (CRC) or other error-correcting or error-detecting information, either in the header, or appended to the data. Referring back to Figure 1, it will be noted that the header for the short packets and for the long packets is not necessarily of the same length, nor is it necessarily in the same format.

The above described PADs 40, 50 are dedicated to either short or long packets. This may offer some advantages in terms of simplicity of hardware required.

In this embodiment, a network also includes a combined PAD 60, which incorporates a data type identifier, with a connection 62 for receiving combined data from a further telephone 42, a further computer 52, and a facsimile machine 64. The PAD 60 identifies the data, and in the case of voice data, assembles the data into short packets, but assembles data from the facsimile machine or computer 52 into long packets.

The three PADs 40, 50, 60 are connected to a multiplexer 70, which combines all packets, both short and long, for transmission along a common transmission medium 72, for example optical fibre, for example using a SONET protocol. It will be appreciated that the multiplexer also serves as a demultiplexer, and performs routing of received data packets to the appropriate PAD. In this embodiment, transmission along the common transmission medium 72 is via separate physical optical fibres or received in transmitted data, as is conventional with ATM network.

It will be appreciated that the functions of the PAD and multiplexer may be distributed differently. For example, packets may be assembled and

disassembled directly within a telephone 42 or a computer 52.

Referring to Fig. 3, one possible implementation of a switch suitable for use in such a network will be described. This implementation has the advantage that it can readily be implemented with only minor modifications to conventional hardware. In the switch 80, an input stream of data is fed to a packet splitter 82. This identifies the incoming data cells as either short or long based on the contents of the additional type flag, strips the type flag from the packets, and produces separated streams of either short or long cells. In this embodiment, the short cells are identical to standard ATM cells, and the cells can be fed directly to a conventional ATM switch. In the simplified embodiment depicted, the switch has a single input, and two outputs, but, in a practical implementation, a complex multichannel switch, such as a General DataComm Apex switch or the like may be employed.

The long packets are fed to a "long ATM" switch. This behaves in a similar manner to a standard ATM switch, but is arranged to handle ATM-type packets having a longer data payload, for example of 1K bytes. The long packets are routed in exactly the same manner as the short packets. Packets coming out of a first set of outputs of the ATM switch 84 and the "long ATM" switch 86 are combined in a first re-packetizer 88a, to provide a first combined output of the switch. In the re-packetizer 88a, a type bit is added to each packet output from the ATM switch and the long ATM switch to regenerate packets of similar format to the input packets. In a similar manner, the second output of the ATM switch and the long ATM switch feed the second re-packetizer 88b, which produces a second output. As will be appreciated further re-packetizers may be connected to further outputs of each of the switches. Likewise, for more than one input, each input would have an associated packet splitter.

An advantage of this arrangement is that the throughput capabilities of the ATM switch 84 and the long ATM switch 86 need not be the same. For

example, in a network in which predominately voice traffic is expected, the ATM switch may be a high bandwidth switch, and the long ATM switch may have a relatively low bandwidth. Then, if the amount of data traffic on the network increases, the long ATM switch 86 can be upgraded, without needing to upgrade the ATM switch 84. Of course, the packet splitter 82 and re-packetizers 88 must be capable of handling the combined throughput for each input/output.

Another advantage with this arrangement is that, the environment seen by each ATM switch 84 is not distinguishable from a standard ATM environment. Indeed, the ATM switch can have inputs or outputs connected directly to a conventional ATM network, in addition to the inputs and outputs which pass through packet splitters and re-packetizers connected to the combined network. This may allow easy upgrading of a existing network, in which the added data-carrying functionality is "bolted on" to an existing ATM network, to enhance the data-carrying capacity of the network without rendering the existing network redundant. Similarly, dedicated data networks may be coupled directly to the long ATM switch 86. This may be useful, for example for connecting a large scale "local" area network directly into the combined network.

It has been previously mentioned that the long and short packets need not necessarily be of similar format. In an alternative implementation, the long data packets may be of similar format to Ethernet data packets, and the packets output by the packet splitter 82 may be of standard Ethernet packet format. In such an embodiment, the switch 86 may be based on a conventional Ethernet router, and may be connected directly into one or more Ethernet local area networks.

It will be appreciated that the above embodiments are purely exemplary, and a number of variations may be made.

For example, if two flag bits are used to signify packet type, this allows four different types of data packet to be sent. One type could be optimised for voice, another optimised for, for example compressed video (for example MPEG-2), another for computer data file transfer, and a fourth for interactive computer data. Alternatively, three fixed length data types and a fourth variable data type could be used, to facilitate transfer of long data blocks. The variable length data packet would, of course, include a length field in the packet header. As mentioned above, a potential drawback of using variable length data packets is that more complicated hardware may be required to perform hardware routing.

An alternative possibility would be to have a slightly larger number of discrete fixed packet lengths. For example, three bits could encode eight different discrete lengths, and the precise lengths could be arbitrary if a look-up table were used to correlate the data type identifier bits to the packet length. If look-up tables are used to correlate packet length to the indicator bits, the look-up tables could be soft-coded, and could be updated dynamically in response to changes in network requirements. For example, traffic analysis software may be used to determine optimum packet sizes based on network usage history, and sizes could be adjusted by sending housekeeping cells to all devices connected to the network to effect an update in packet size. In such an arrangement, it would be preferable to have indicator bits capable of encoding more packet lengths than are used at any one time on the network. For example three bits may be used to encode packet size, but only four different packet sizes may be used at any one time. Assuming that packets codes 0-3 are defined, housekeeping cells may be issued to all devices indicating that henceforth packet code 4 corresponds to a new desired packet size. Thereafter, a command may be issued prohibiting further use of packet code 0. Thus, data in transit using the packet size corresponding to packet code 0 may be successfully delivered, but thereafter, a network will use packet size corresponding to packet code 4 in place of the packet size corresponding to packet code 0. This determination of packet size may be made, as discussed, by network management apparatus, based on analysis of network

5 traffic. It is also possible for a user who expects to transmit a significant amount of data which can conveniently be packaged into packets of a known size to request definition of an appropriate size packet. Usually, however, the amount of network bandwidth consumed by an individual user will not be sufficient to justify re-configuring the network for an individual request, so it is not desirable that a user request controls network configuration directly. User requests may, however, be used by the network controlling software in determination of appropriate packet sizes for the network.

10 To facilitate inter-connection with other networks, it may be desirable to provide means for converting a long packet into a series of shorter packets, or vice versa. Particularly where the embodiment uses an extended ATM-based arrangement, as described above, it is desirable to provide means for converting a long data packet into a plurality of ATM cells. In this way, a large network including a conventional ATM network and a network according to the
15 embodiment may be freely mixed, and, at the boundaries of the network according to the embodiment, data may continue to be routed to its destination over the conventional ATM network.

Referring now to Fig. 4, a second embodiment will now be described; this schematic diagram illustrates the order of data transmission when a high
20 priority packet interrupts a lower priority packet. A first, lengthy, data packet is transmitted as a header 100 including a priority flag 101 and an initial data portion 102a. During transmission of the initial data portion 102a, a higher priority packet is received for sending along the same path. The higher priority packet comprises a header portion 120 including a priority flag 121, in this
25 case set to high priority, and a data portion 122. To send this packet without having to wait until the first packet has finished, an interrupt code 110 is inserted in the data portion 102a, which signals to an appropriately configured receiver that sending of the first packet has been interrupted. The higher priority packet is then sent, followed by a resumption code 111, followed by the
30 remainder 102b of the data of the first packet.

Referring to Fig. 5, apparatus for performing the interruption described above will now be described. A network interface 130 is arranged to receive data packets for onward transmission. Received packets pass by priority detector 132 which determines whether interruption is required, by reference to the priority of the packet in transit stored in current priority store 133. In the normal course of events, packets are passed to transmitter 134 in the order received for transmission over a network, here a 1Gbit/s ethernet 138. However, if the priority detector determines that interruption is required, a signal is sent to interrupt device 136 which inserts an interrupt code, here chosen to be a pre-determined "spare" code from the 10bit coding scheme used for the Ethernet, buffers the remainder of the packet in transit as necessary, and instructs the transmitter to commence output of the higher priority packet. After the higher priority packet has been sent, a resume code is sent (this may not always be necessary, as the header 120 of the higher priority packet will contain the packet length, but this may simplify reassembly of the fragmented packets and increase fault tolerance), followed by the remainder of the first packet.

Signalling of interruption is preferably implemented at the physical layer, for example by using an otherwise unused 10 bit code in the case of Ethernet.

It will be appreciated from the above that the precise details of the implementation may be varied considerably, and also that each of the features discussed above may be provided independently in the context of other embodiments, unless otherwise stated.

Claims

1. A method of communicating data over a network comprising transmitting a first data packet, receiving a higher priority data packet, interrupting transmitting the first data packet to transmit the higher priority data packet, and
5 resuming transmission of the first data packet.

2. A method according to Claim 1, wherein interrupting includes inserting a code signifying interruption into the transmitted data packet.

3. A method according to Claim 1 or Claim 2, further comprising maintaining a count of interruptions to a lower priority packet and declining
10 interruption by a higher priority packet when a maximum count is reached.

4. A method according to any preceding claim wherein interruption is signalled at the physical layer.

5. A method according to any preceding claim wherein only two priority levels are provided.

6. A method according to any of Claims 1 to 4, wherein multiple priority
15 levels are provided.

7. Apparatus for communicating data over a network comprising means for transmitting a first data packet, means for receiving a higher priority data packet, means for interrupting transmitting the first data packet to transmit the
20 higher priority data packet, and means for resuming transmission of the first data packet.

8. Apparatus according to Claim 7 including means for signalling interruption at the physical layer.

9. A method of communicating data over a network comprising selecting
25 one of at least two predetermined packet sizes based on the data type and

transmitting a packet of the selected size over the network.

10. A method according to Claim 9, wherein the packet formats differ depending on the data type.

5 11. A data network including means for transmitting data of a first type as a series of packets of a first length over a communication medium and means for transmitting data of a second type as packets of a second length over the same communication medium.

10 12. A network according to Claim 11 further including routing means connected to the communication medium for routing packets of both said first length and said second length.

15 13. A transmitter for connection to a communication network comprising means for selecting a data packet format or length from a plurality of predetermined formats or packet lengths; means for providing a header for the data packet, the header including at least one flag bit indicating the length or format of the data packet; and means for transmitting the data packet over the communication medium.

20 14. A receiver for receiving data from a communication network, the receiver including means for reading a flag bit contained in the header of a data packet transmitted on the communication medium; means for determining the size or format of data packet to receive based on the information; and means for receiving a data packet of said size or format.

15. A network or network device substantially as any one herein described or as illustrated in any figure.

25 16. A method of communicating data substantially as any one herein described.



Application No: GB 9821555.1
Claims searched: 1-8

Examiner: Gareth Griffiths
Date of search: 17 March 2000

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK CI (Ed.R): H4K (KTK), H4P (PPNB, PPS)

Int CI (Ed.7): H04L 12/56

Other: Online Databases: WPI, EPODOC, JAPIO

Documents considered to be relevant:

| Category | Identity of document and relevant passage | Relevant to claims |
|----------|--|--------------------|
| X | GB2187917 A (UNIVERSITY OF STRATHCLYDE) p.3 line 126 - p.4 line 69 | 1, 7 at least |
| X | GB2166320 A (STC) whole document | 1, 7 at least |
| X | WO95/12265 A1 (NORTHERN TELECOM) p.7 lines 3 - 11 | 1, 7 at least |
| X | US5251209 (FOUDRIAT) col.8 lines 6 - 19 | 1, 7 at least |
| X | JP580173932 A (OKI) see abstract | 1, 7 at least |

| | | | |
|---|---|---|--|
| X | Document indicating lack of novelty or inventive step | A | Document indicating technological background and/or state of the art. |
| Y | Document indicating lack of inventive step if combined with one or more other documents of same category. | P | Document published on or after the declared priority date but before the filing date of this invention. |
| & | Member of the same patent family | E | Patent document published on or after, but with priority date earlier than, the filing date of this application. |